arXiv:1511.01946v1 [cs.AR] 5 Nov 2015

SecureD: A Secure Dual Core Embedded Processor

Roshan G. Ragel[1], Jude A. Ambrose[2], and Sri Parameswaran[2]

[1]Department of Computer Engineering, University of Peradeniya 20400 Sri Lanka
[2]School of Computer Science and Engineering, University of New South Wales NSW 2052 Sydney Australia

Abstract

Security of embedded computing systems is becoming of paramount concern as these devices become more ubiquitous, contain personal information and are increasingly used for financial transactions. Security attacks targeting embedded systems illegally gain access to the information in these devices or destroy information. The two most common types of attacks embedded systems encounter are code-injection and power analysis attacks. In the past, a number of countermeasures, both hardware- and software-based, were proposed individually against these two types of attacks. However, no single system exists to counter both of these two prominent attacks in a processor based embedded system. Therefore, this paper, for the first time, proposes a hardware/software based countermeasure against both code-injection attacks and power analysis based side-channel attacks in a dual core embedded system. The proposed processor, named *SecureD*, has an area overhead of just 3.80% and an average runtime increase of 20.0% when compared to a standard dual processing system. The overhead were measured using a set of industry standard application benchmarks, with two encryption and five other programs.

Keywords

Side Channel Attacks Code Injection Attacks Hardware/Software Co-design

1 Introduction

Multiple processors have been used on a single chip as the increasing functionality of embedded systems demands more processing power. Consumer devices such as mobile phones already deploy dual processors in their designs [46]. An increasing number of such embedded systems have to overcome security threats. Potential targets

of adversaries vary from low-end systems such as wireless handsets, networked sensors and smart cards, to high-end systems such as network routers, gateways, firewalls and servers.

Security threats in embedded systems could be classified by the means used to launch attacks. Typical launch methods are: physical, logical and side-channel. Physical attacks refer to unauthorized physical access to the embedded system itself and are feasible only when the attacker has direct access to the system. Such attacks are identified by packaging mechanisms, such as tamper evident packing (make an attempted attack apparent so that subsequent inspection will show an attack had been attempted) [45]. Logical attacks exploit weaknesses in logical systems such as software or a cryptographic protocol to gain access to unauthorized information. Logical attacks are deployed easily against systems which are able to download and execute software and have vulnerabilities in their design. Side-channel attacks are performed by observing properties of the system (such as power consumption or electromagnetic emission) while the system performs cryptographic operations. Unlike physical attacks, logical and side-channel attacks are application specific and could be diagnosed and prevented by software and/or architectural techniques.

Most recent logical attacks result in demolishing code integrity of an application program [31]. They dynamically change instructions with the intention of gaining control over the execution flow of a program. Attacks that are involved in violating software/code integrity are called code-injection attacks. Code-injection attacks often exploit common implementation mistakes in application programs, which are referred to as security vulnerabilities. The number of malicious attacks always increases with the amount of software code [11]. The most known side-channel attack is the power analysis attack [24], where secret keys used in an encryption program were successfully predicted by observing the power dissipation from a chip. Devices like *smart cards* [10, 12], *PDAs* [17] and *mobile phones* [46] have microprocessor chips built inside, performing secure transactions using secret keys. Power dissipation/consumption of a chip is the most exploited property used to predict secret keys using side channel attacks [26, 48].

A number of hardware- and software-based techniques, such as ARM® *TrustZone* [3] and *ARMOR* [2] respectively, were proposed in the past as secure infrastructures. However, they are not designed to detect/prevent the most frequently encountered logical and side-channel attacks on embedded systems: code-injection and power analysis attacks. There exist a number of individual detection mechanisms for these two attacks in both software and architecture domains. However, to our knowledge, no single system was proposed that is capable of thwarting or warning against both these attacks. Therefore, this paper, *for the first time*, presents a secure dual core embedded system, we call *SecureD* (for **Secure D**ual core), which both detects code-injection attacks and protects against power analysis attacks. *SecureD* considers security as one of its design objectives (as opposed to security being an afterthought in the design) [25] and uses a hardware/software solution with microarchitectural changes to the processor design, which makes the integration of security into the design and implementation much faster and easier, while considerably reducing the overhead.

Paper Organization

The rest of the paper is organized as follows. Section 2 summarizes previously proposed countermeasures against both code-injection and power analysis attacks. The microarchitecture description of *SecureD* is provided in Section 3. Section 4 illustrates the design flow of our system. The experimental setup is presented in Section 5 and the results are presented in Section 6. Finally, the paper is concluded in Section 8.

2 Related Work

In this section, we present a range of prior detection and prevention techniques, both hardware- and software-based against both code-injection and power analysis attacks.

A wide range of techniques have been suggested in the past to detect/counter code-injection attacks. They could broadly be categorized into software based and hardware assisted techniques. Software-based techniques use software tools and methods to overcome these attacks without changing the microarchitecture of the processor. Hardware assisted techniques use additional hardware blocks or microarchitectural support to detect code-injection attacks. Software-based techniques can be further categorized into two: static and dynamic. Static techniques try to detect vulnerabilities at compile time. Wagner et al. propose an automated static code analysis tool to detect code that might invite buffer overflow attacks [44], but this produces a relatively large number of false positives. Another static technique uses a language that has only the safe constructs of another language, for example, Cyclone [23]. Dynamic software based techniques avoid or considerably reduce code injection attacks at runtime and they either use formal methods to prove a program behaves as expected or use software constructs to monitor proper program behaviour at runtime. *Proof-Carrying Code* [33] is an example of the former and *Stack Guard* [14] the latter.

Hardware assisted techniques are mainly attack-specific. A number of researchers [27, 29, 47] propose architectural detection of buffer overflow attacks which are a type of code-injection attacks. Code-injection attacks can be detected by monitoring code integrity of a program at runtime. Arora et al. [8] use an additional co-processor and hardware tables to perform software/code integrity checks. This system identifies program properties at different levels of granularity and stores multiple control flow levels of data and checksums to perform software integrity monitoring. This method produces code that is not relocatable due to the pre-generated hardware tables. Ragel and Parameswaran in their secure processor called *IMPRES* [35] use similar program properties to that of [8], but only perform check-summing at the basic block level and therefore reduce the complexity of their solution. However, *IMPRES* only detects code-injection attacks and is designed for a single processor embedded system [36].

Similar to detection techniques for code-injection attacks, countermeasures against power analysis attacks can be classified into software-based and hardware-based. Masking and current flattening are the two major software-based countermeasures. Table and data masking techniques [13, 16, 19, 30] use random values during the actual computation to prevent the processed data being exploited by the adversary. Muresan and Gebotys [32] proposed a current flattening technique, where the dissipated

current is flattened by adding *no-op*s in the code to provide sufficient discharge. Ambrose et al. [6] proposed a randomized instruction injection technique, where dummy instructions were injected during the actual execution. Even though this technique has been proven effective for small number of data samples, the phase substituition techniques [18] can be used to isolate the injected effects with large number of samples. Authors in [43] present a comprehensive study on shuffling, which is similar to masking. It was mentioned in [43] that the shuffling is effective when both the execution order and the physical resource usage are randomized.

Non-deterministic processing and hardware balancing for power are the most explored hardware-based countermeasures against power analysis attacks. *Non-Deterministic Processors* [28] execute instructions out-of-order, issuing independent code segments randomly during runtime and therefore preventing the adversary from identifying the places where specific instructions are executed. *Dual-Rail* circuits [15] (or *Dual-Rail Pre-charge* (DRP) logic [34]) are designed to consume constant power regardless of data processed. In Dual-Rail circuits, each logic circuit is attached to another similar logic circuit, complementing the discharge occurring in the original logic circuit due to bit-flips [37]. Other similar hardware balancing techniques are presented in [21, 34, 41, 42]. Residue Number Systems (RNS) based randomized hardware approach is presented in [5], randomly chooses the moduli sets to obfuscate the binary double and add operation from the power profile. This technique requires large data inputs to gain advantage in performance, since RNS circuits cost in hardware. Furthermore, the proposed technique in [5] is not complete and the authors have not tested the point double and add for ECC. Tanimura and Dutt [38] propose an improvement for WDDL [42] using a simulated annealing approach to automatically generate complementary cells to reduce the information leakage from the power profile. This circuit level balancing ExCCel approach [38] was further improved in [39, 40].

MUTE [7], a hardware/software information balancing technique, is a system level microarchitectural countermeasure to protect against power analysis attacks. *MUTE* uses two processors to execute the same encryption program with the actual and complemented data, balancing information (against power). *MUTE* requires minimal software instrumentation compared to the current flattening technique in [32]. It utilizes an identical second core of the dual core processor and therefore requires minimal hardware (for additional control logic) compared to the hardware balancing methods [15, 21, 22, 34, 37, 41, 42]. However, *MUTE* protects the system only against power analysis attacks.

The dual core secure embedded processor system presented in this paper, *SecureD*, makes use of techniques similar to that of *IMPRES* [35] to detect code-injection attacks and to that of *MUTE* [7] to protect against power analysis attacks. As far as we are aware, there has been no embedded processor system, which is immune to both code-injection and power analysis attacks.

2.1 Contributions

- A secure dual core embedded system is proposed to both detect code-injection attacks and protect against power analysis attacks. As far as we are aware, this is the first time an MPSoC solution is proposed to safeguard against both power

analysis and code injection, which are two of the major security threats in embedded systems.

- The processor is capable of detecting bit-flips including those causing control flow errors with minimal latency.
- An interrupt handling mechanism for an ASIP MPSoC is presented.
- A novel switching and synchronizing mechanism for information balancing to counter power analysis is proposed which was not available in [7].
- A rapid simulation and synthesis environment is used for a dual core processor.

2.2 Limitations and Assumptions

- Our technique addresses only multiprocessor embedded systems with at least two identical (homogeneous) processors. For heterogeneous multiprocessor systems, identical functional units for encryption/decryption has to be used in two of the processors.
- We assume that our system is self contained with separate memories for each of the processors.
- Both processors are clocked by a single source.

3 Processor Architecture

In this section, we give an overview of *SecureD* architecture. We further discuss how the code-injection detection technique similar to the one presented in [35] and the power analysis protection similar to the one presented in [7] are integrated into *SecureD*.

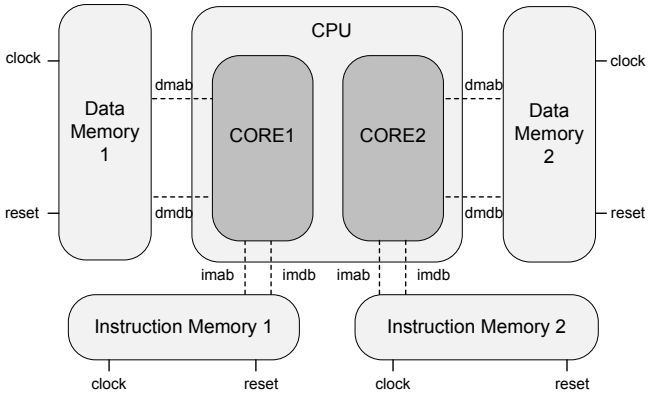

Figure 1: The Base (Non Secure) Dual Core Processor with Memory Modules

Figure 1 depicts the schematic diagram of the base dual core processor taken for our design. As depicted, the processor has two identical cores with separate instruction and data memories for each core. The base processor is statically scheduled to run different threads/applications on each of its cores.

Figure 2 depicts how one of the cores from the base processor is altered so that it is capable of detecting code-injection attacks and preventing power analysis attacks.

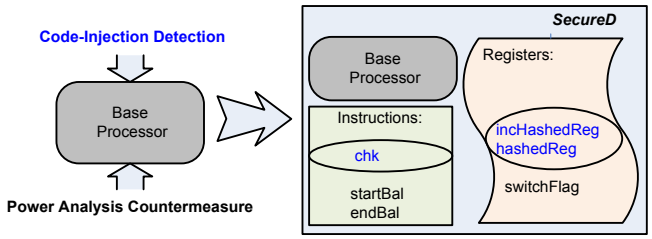

Figure 2: Design Alterations to One of the Cores of the Base Processor

As depicted in Figure 2, each core of our system will be augmented by adding registers and logic to handle additional instructions. The additional instructions and registers encircled (see Figure 2) are used to detect code-injection attacks and the rest to prevent power analysis attacks. Following subsections will detail architectural modifications and how they will be used in ensuring security in our *SecureD* processor.

3.1 *SecureD* for Detecting Code-injection Attacks

SecureD ensures code integrity of an application by verifying whether all basic blocks of the program are intact at runtime [35]. This is achieved by performing basic block integrity checks at runtime. The basic block integrity checker incorporates the following tasks: (1) identifying basic blocks and calculating and assigning checksums for each basic block at compile time; and (2) re-calculating the checksums at runtime and comparing them with loaded static values.

At compile time, an application program is grouped into basic blocks based on the control flow of the application. Then, each basic block is processed separately to calculate a checksum based on the instructions of that block. The calculated checksum is then inserted at the beginning of each basic block using a special instruction (*chk* instruction).

At runtime, the first instruction of a loaded basic block is a *chk* instruction that carries the checksum for the corresponding basic block. When an instruction of this kind is fetched, the checksum is loaded into a special register (*hashedReg*). The checksum for each basic block is incrementally re-calculated at runtime while instructions belonging to the basic block are executed and is stored in another special register (*incHashedReg*). The last instruction of each basic block is a control flow instruction (CFI), and if not, one is inserted at compile time (if it is not present at the end of a basic block). CFIs are altered such that they will (a) incrementally store checksums of their own, and (b) compare the result against the one loaded through *chk* instructions. A mismatch in the comparison will indicate a code integrity violation and generate an exception.

3.2 *SecureD* for Preventing Power Analysis Attacks

When there is an encryption program to be executed, *SecureD* will use both of its cores, where the first one will execute the original encryption program, while the second core executes the complementary program in parallel. Similar to [7], the original and complementary programs are devised by having the same instruction sequence but complementary data processing. Hence, a clock cycle accurate synchronization is preserved, masking the information signature in the power profile. The memories are setup statically, making the processor aware of where the programs are stored and which data locations are used for different programs. The term *balancing* is used for this process as was named in [7].

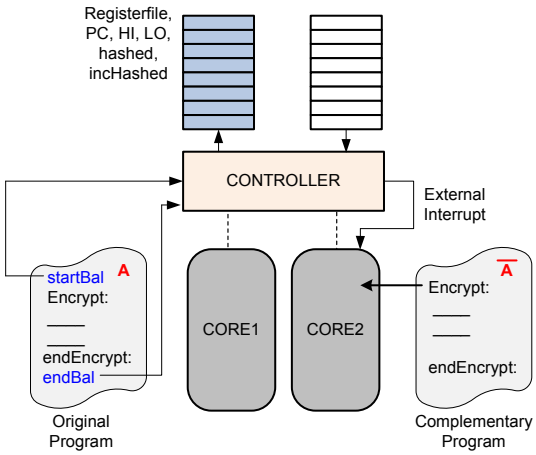

Figure 3: Switching and Synchronizing

The *balancing* is triggered and terminated by two special instructions which are instrumented in the source code. As depicted in Figure 3, the execution of the *startBal* instruction indicates the *CONTROLLER* that an encryption program is scheduled in the core. An External Interrupt is sent to *CORE2*. This is a maskable interrupt which will be triggered after all the pipelines are flushed. Necessary registers (such as the register-file and program counter) of *CORE2* are saved in the stack as shown in Figure 3. After the registers are saved, the *CONTROLLER* sends the program counter (PC) values to *CORE1* and *CORE2* on the same clock cycle (i.e., PC values of Encrypt in program A and $\overline{A}$). Both the original and complementary programs are executed in parallel by *CORE1* and *CORE2* respectively. When the encryption is completed, the *endBal* instruction in program A will be executed by *CORE1*, which will send a signal to the *CONTROLLER* indicating the completion of encryption. The *CONTROLLER* restores the saved registers from the stack. *CORE2* will resume its execution and *CORE1* will be scheduled with the next program to be executed.

3.3 Interrupt Handling and *SecureD*

This subsection describes how interrupts are handled by *SecureD* in two different scenarios: one, when *SecureD* is executing a regular application and two, while *SecureD* is executing an encryption algorithm (that is during *balancing*). For both the scenarios, interrupts will be serviced after the pipelines are flushed. Hence, the registers are updated with correct values. Each interrupt routine will have three code segments: one, instructions to save all the registers into the stack; two, the actual code for the routine; and three, instructions to restore the registers from the stack and the end instruction to terminate the interrupt. The end interrupt instruction sends a non-maskable interrupt (NMI) to the *CONTROLLER*, which will force the controller to change the PC to resume the original execution. In regular interrupt handling, the *CONTROLLER* will change only the PC of the core which was servicing the interrupt and lets the other core perform its normal execution.

However, handling an interrupt while the core is performing *balancing* is slightly different. When a core receives an interrupt during balancing, the other core is put on hold until the interrupt is serviced. The core on hold can also be allowed to execute the next task in the queue, but with careful modification in the *CONTROLLER* to maintain synchronization for balancing. The NMI request will force the *CONTROLLER* to change the PC of both cores to their original location to resume balancing. This is done on the same clock cycle in-order to preserve synchronization.

4 Design Flow

In this section, an overview of the proposed design flow for the *SecureD* architecture is discussed. First, the design of a software interface that allows the application to interact with the architectural enhancement is described, and then the design of the architectural enhancement itself is discussed.

4.1 Software Design Flow

Figure 4 depicts the software design flow of our framework. The source code of the application program in C is compiled with the front-end of a compiler to generate the assembly version of the application (*.s files*). Special instructions for both code-injection detection (*chk* instructions with checksum values of basic blocks) and power analysis protection (*startBal* and *endBal*) are instrumented into the assembly. The resulting assembly files are assembled and linked to generate the binary (Instrumented Binary in Figure 4) of the application.

4.2 Hardware Design Flow

Figure 5 depicts the generation of a processor model that implements *SecureD*. The Instruction Set Architecture (ISA) is fed into an automatic processor design tool (*ASIP-Meister* [1]) to generate each core of the processor model. Specifications of special instructions for *SecureD* are given to *ASIPMeister*. Additional registers are chosen to

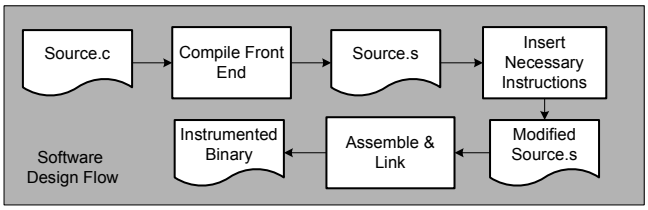

Figure 4: Software Design Flow

implement secure functionalities. Interrupts are selected and implemented for switching and synchronizing processors. The logical functionalities of the special instructions and interrupts are specified as micro-instructions. The single core processor model in VHDL is generated using the *generate hardware module* of *ASIPMeister*.

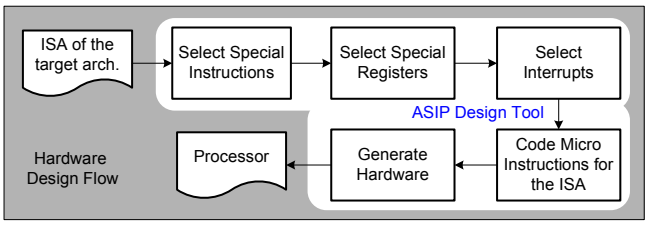

Figure 5: Hardware Design Flow

ASIPMeister generates a synthesizable VHDL description of a single processor model, thus two identical processors are generated separately and are combined manually to make the VHDL model of *SecureD*. Additional components such as the *CONTROLLER* as shown in Figure 3 are added during the manual combining process.

5 Experimental Setup

Our *SecureD* framework is implemented using processors with the PISA (*Portable Instruction Set Architecture*) instruction set (as implemented in *SimpleScalar Tool Set* with a six stage pipeline) processor without cache. Figure 6 illustrates the experimental setup, identifying key elements and tools. Programs in C are compiled using GNU/GCC cross compiler for the PISA instruction set. The *Instrumented Binary* is produced as explained in Section 4.1. The processor model represents the *SecureD* processor in VHDL description.

Synopsys® *Design Compiler* is used to synthesize the processor model as shown in Figure 6. Hardware details are reported by the *Design Compiler*. The instrumented binary is processed to set up both data and program memories. Modelsim VHDL simulator is used to simulate the synthesized processor and the memory models to deduce the runtime (in number of clock cycles).

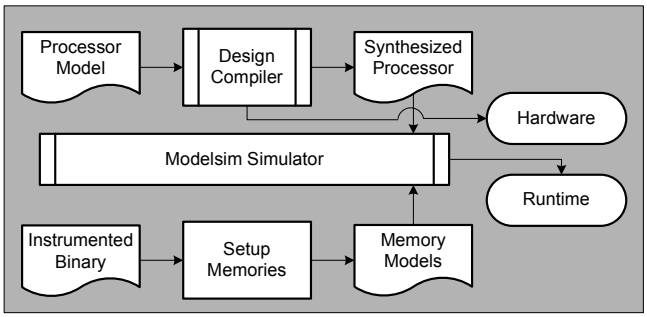

Figure 6: Experimental Setup

6 Results

This section presents the hardware and runtime overhead details of our *SecureD* processor, comparing with the processors implemented for only detecting code-injection attacks - *SecureD-I* [35], and for only protecting against power analysis attacks - *SecureD-M* [7].

6.1 Hardware Summary

Table 1 tabulates the hardware summary produced by Design Compiler for all processors investigated. The CMOS 65nm technology is used. *Non-SecureD* is a processor with dual cores and without any secure implementations. *SecureD-M* costs an additional 1.9% and 2.7% in gates and cells respectively, while *SecureD-I* costs an additional 1.3% and 1.5% compared to the *Non-Secure* processor. *SecureD* costs 3.8% area overhead in both gates and cells.

Table 1: Hardware Details

	Area (# of gates)	**Cell Area** (# of cells)	**Clock** (ns)
Non-SecureD	484180	132828	31.58
SecureD-M	493330	136496	31.06
SecureD-I	490715	134812	31.94
SecureD	502830	137906	31.80

The *SecureD-M* processor has a slight decrease in the clock width of the critical path because of the optimizations performed by the Design Compiler. *SecureD-I* has a slight increase in the clock due to its runtime checksumming logic. *SecureD* expectedly has an increased clock width compared to *Non-SecureD* and *SecureD-M* and a slightly decreased clock width compared to *SecureD-I* because of the optimizations.

Hardware Summary in FPGA

The base dual core and *SecureD* processors were implemented on an FPGA (XC2V3000-4FG676) and the hardware details are presented in Table 2. *SecureD* consumes an additional 10.8% in slices, 2.8% in LUTS, 1.5% in IOBs compared to the non-secure dual core processor *Non-SecureD*. Clock width has slightly increased in addition to the hardware overhead in logics, such as LUTs, IOBs and slices.

Table 2: Hardware Details in XC2V3000-4FG676

	Non-SecureD	***SecureD***
Hardware		
Slice Flip-flops (28672)	4300	4823
4 input LUTs (28672)	16476	16942
Bonded IOBs (842)	410	416
GCLKs (16)	1	1
Timing		
Clock (ns)	209	216
Logic Levels	515	529

6.2 Runtime Analysis

Figure 7 depicts the runtime analysis of all four configurations of our dual core processors (*Non-SecureD*, *SecureD-M*, *SecureD-I* and *SecureD*) for different benchmark applications. Out of the seven applications tabulated, two (DES and AES) are considered to be encryption algorithms and *balancing* is enabled while they are executed. Each cell in the table (see Figure 7) has 4 parts, where the top left part specifies the runtime of a particular application on the non-secure dual core (*Non-SecureD*), the top right part represents the runtime on *SecureD-M* processor, the bottom left part indicates the runtime on *Secured-I* processor and the bottom right part has the runtime of an application on *SecureD* processor. The *Non-SecureD* neither includes the balancing feature nor the code injection detection feature, whereas the *SecureD-I* only includes the code injection detection feature. The *SecureD-M* design only contains the balancing solution to prevent power analysis whereas the *SecureD* processor prevents both the attacks.

When two applications are scheduled on dual cores, the runtime of the application which is higher is considered for the net runtime of the two applications. Hence, the cells along the diagonal of the table (in Figure 7) represent the runtime of each application alone. Since *adpcm.encode* has the highest runtime amongst all programs, combinations with *adpcm.encode* will be having the runtime of *adpcm.encode* (e.g., see *Non-SecureD* cells in column two of Figure 7).

In *SecureD-M* and *SecureD* configurations, when either AES or DES (encryption algorithms) is scheduled to one of the cores, the application scheduled on the second core has to wait until the *balancing* of AES or DES comes to an end (i.e., when *balancing* is performed, both cores will be executing the same program, either AES or

DES, with complementary values). Note that, currently, *balancing* is performed only for AES and DES in the given set of benchmark applications. For example, the runtime of applications *adpcm.encode* and *DES encrypt* scheduled on *SecureD-M* (see Figure 7 row 2, column 2), is the addition of the runtime of *DES encrypt* on both cores (for *balancing*), the time taken for switching, and the runtime of *adpcm.encode* on one of the cores which is performed after the *balancing* finishes.

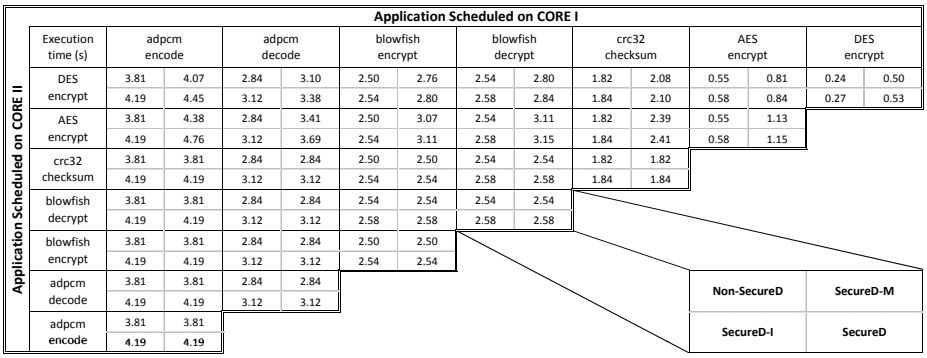

Figure 7: Runtime Comparisons

As depicted in Figure 7, the cells corresponding to *SecureD-I* consume additional time due to the runtime check-summing used for detecting code integrity violations. The average runtime overhead of *SecureD-I* for all seven applications is 6.06%. Compared to the runtime of *Non-SecureD* configuration, *SecureD-M* consumes extra runtime only when either AES or DES is scheduled on one of or both the cores. The average runtime overhead of all seven applications on *SecureD-M* is 15.4% compared to *Non-SecureD*. As *SecureD* integrates both code-injection detection and power analysis protection, it consumes extra runtime similar to that of *SecureD-I* when no application with encryption (such as AES or DES) is present. *SecureD* consumes additional runtime compared to *SecureD-I* when applications such as AES or DES are scheduled to one or both of the cores. The average runtime overhead of *SecureD* with two encryption and five regular applications is 20.0%.

Table 3 presents the delay in number of clock cycles for switching and interrupt servicing during balancing. All 37 registers including the 32 registers in the register-file are saved and restored sequentially. Every time an external interrupt is fired (to switch to balancing or to service interrupt) there is a 6 clock cycle delay to flush the pipeline to update all the registers. A clock cycle delay is consumed for each interrupt call or switch instruction (*startBal*). Another clock cycle delay is needed to exit the interrupt or to exit *balancing* (using *endBal* instruction).

Table 3: Switching and Interrupt Servicing Delay

	Clock Cycles
Delay	
Store Register-file	320
Store PC,HI,LO registers	30
Store incHashed, hashed	20
Restore Register-file	320
Restore PC,HI,LO registers	30
Restore incHashed, hashed	20
Flush Pipelines	6
Interrupt to switch	1
Exit the interrupt	1
Total Delay	748

7 Discussions

This paper focuses on power analysis based side channel attacks. Our solution can be further applied to electromagnetic based side channel attacks [20], however not tested. Previous techniques have proven that the DPA resistant techniques have also prevented DPA, since electro-magnetic emmissions are closely related to power dissipations [16, 18].

We assume that the processors are designed and layed out as homogeneous components which will dissipate around the same amount of power. However, this could be a challenging task and would require manual design considerations. We could use the CoRaS solution [4] to counteract this limitation by randomly swapping rounds of the encryption/decryption between processors.

We assume that we do not have caches in our MPSoC or disable caches during the balancing operation. This is to make sure we execute both the processors in instruction-level lock-step mode. Cache locking mechanisms [9] can be utilized during secure execution, if we want to implement cache to enhance performance.

8 Conclusions

In this paper, we have presented *SecureD*, a dual core based multiprocessor, which is immune to code-injection attacks and power analysis based side-channel attacks. Both these attacks are considered to be the most successful attacks in embedded systems in recent years. *SecureD* protects against code-injection by performing static time basic block instrumentation (with checksums) and runtime comparison for detecting code integrity violations and power analysis is prevented by implementing information balancing to mask the actual information (or the secret key) from the power profile. *SecureD* is easily scalable beyond dual processors and has minimal performance and hardware overhead compared to previous techniques proposed to separately handle the attacks.

References

[1] The PEAS Team. ASIP Meister, 2002. Available at: http://www.asip-solutions.com/english/.

[2] ARMORs: A Software-Implemented Fault Tolerance Environment, 2007. Available at: http://www.crhc.uiuc.edu/DEPEND/projects/ARMOR/ARMOR.html.

[3] TrustZone: Security Foundation by ARM, 2008. Available at: http://www.arm.com/products/esd/trustzone_home.html.

[4] J. Ambrose, A. Ignjatovic, and S. Parameswaran. CoRaS: A multiprocessor key corruption and random round swapping for power analysis side channel attacks: A DES case study. In *ISCAS*, pages 253–256, 2012.

[5] J. A. Ambrose, H. Pettenghi, and L. Sousa. Darns: A randomized multi-modulo rns architecture for double-and-add in ecc to prevent power analysis side channel attacks. In *ASP-DAC*, 2013.

[6] J. A. Ambrose, R. G. Ragel, and S. Parameswaran. Randomized instruction injection to counter power analysis attacks. *ACM Trans. Embed. Comput. Syst.*, 2012.

[7] J. A. Ambrose, R. G. Ragel, S. Parameswaran, and A. Ignjatovic. Multiprocessor information concealment architecture to prevent power analysis-based side channel attacks. *IET computers & digital techniques*, 5(1):1–15, 2011.

[8] D. Arora, S. Ravi, A. Raghunathan, and N. K. Jha. Secure embedded processing through hardware-assisted runtime monitoring. In *Proceedings of the Design, Automation and Test in Europe (DATE'05)*, volume 1, 2005.

[9] A. Asaduzzaman, F. N. Sibai, and M. Rani. Improving cache locking performance of modern embedded systems via the addition of a miss table at the L2 cache level. *J. Syst. Archit.*, pages 151–162, 2010.

[10] E. Biham and A. Shamir. Power Analysis of the Key Scheduling of the AES Candidates. In *In Second Advanced Encryption Standard (AES) Candidate Conference*, pages 343 – 347, 2003.

[11] CERT Coordination Center. Vulnerability Notes Database, CERT Coordination Center, 2007.

[12] S. Chaumette and D. Sauveron. New security problems raised by open multiapplication smart cards.

[13] J.-S. Coron and L. Goubin. On Boolean and Arithmetic Masking against Differential Power Analysis. In *Ches '00*, pages 231–237, London, UK, 2000.

[14] C. Cowan, C. Pu, D. Maier, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, Q. Zhang, and H. Hinton. StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks. In *Proc. 7th USENIX Security Conference*, pages 63–78, San Antonio, Texas, 1998.

[15] S. Danil, M. Julian, B. Alexander, and Y. Alex. Design and analysis of dual-rail circuits for security applications. *IEEE Trans. Comput.*, 54(4):449–460, 2005.

[16] C. Gebotys. A Table Masking Countermeasure for Low-Energy Secure Embedded Systems. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 14(7):740–753, 2006.

[17] C. H. Gebotys and R. J. Gebotys. Secure Elliptic Curve Implementations: An Analysis of Resistance to Power-Attacks in a DSP Processor. In *CHES '02: Revised Papers from the 4th International Workshop on Cryptographic Hardware and Embedded Systems*, pages 114–128, London, UK, 2003. Springer-Verlag.

[18] C. H. Gebotys and B. A. White. A Phase Substitution Technique for DEMA of Embedded Cryptographic Systems. In *ITNG*, pages 868–869, 2007.

[19] L. Goubin and J. Patarin. DES and Differential Power Analysis (The "Duplication" Method). In *CHES '99: Proceedings of the First International Workshop on Cryptographic Hardware and Embedded Systems*, pages 158–172, London, UK, 1999. Springer-Verlag.

[20] N. Homma, T. Aoki, and A. Satoh. Electromagnetic information leakage for side-channel analysis of cryptographic modules. In *EMC*, pages 97–102, 2010.

[21] D. Hwang, K. Tiri, A. Hodjat, B.-C. Lai, S. Yang, P. Schaumont, and I. Verbauwhede. Aes-Based Security Coprocessor IC in 0.18um CMOS With Resistance to Differential Power Analysis Side-Channel Attacks. *IEEE Journal of Solid-State Circuits*, 41(4):781– 792, 2006.

[22] D. D. Hwang, P. Schaumont, K. Tiri, and I. Verbauwhede. Securing embedded systems. *IEEE Security and Privacy*, 4(2):40–49, 2006.

[23] T. Jim, G. Morrisett, D. Grossman, M. Hicks, J. Cheney, and Y. Wang. Cyclone: A Safe Dialect of C. In *In USENIX annual technical conference, Monterey, CA, June 2002.*, 2002.

[24] P. Kocher, J. Jaffe, and B. Jun. Differential power analysis. *Lecture Notes in Computer Science*, 1666:388–397, 1999.

[25] P. Kocher, R. Lee, G. McGraw, A. Ragunanthan, and S. Ravi. Security as a new dimension in embedded system design, June 2004.

[26] F. Koeune and F.-X. Standaert. A Tutorial on Physical Security and Side-Channel Attacks. In *Foundations of Security Analysis and Design III : FOSAD 2004/2005*, pages 78–108, 2006.

[27] R. Lee, D. Karig, J. McGregor, and Z. Shi. Enlisting hardware architecture to thwart malicious code injection. In *Proceedings of the International Conference on Security in Pervasive Computing*. Springer Verlag LNCS, March 2003.

[28] D. May, H. L. Muller, and N. P. Smart. Non-deterministic Processors. In *ACISP '01: Proceedings of the 6th Australasian Conference on Information Security and Privacy*, pages 115–129, London, UK, 2001. Springer-Verlag.

[29] J. McGregor, D. Karig, Z. Shi, and R. Lee. A processor architecture defense against buffer overflow attacks. In *Proceedings of the International Conference on Security in Pervasive Computing (SPC-2003)*, pages 237–252. Springer Verlag, March 2003.

[30] T. S. Messerges, E. A. Dabbish, and R. H. Sloan. Examining smart-card security under the threat of power analysis attacks. *IEEE Trans. Computers*, 51(5):541–552, 2002.

[31] M. Milenkovic, A. Milenkovic, and E. Jovanov. Hardware support for code integrity in embedded processors. In *CASES '05*, pages 55–65, New York, NY, USA, 2005. ACM Press.

[32] R. Muresan and C. H. Gebotys. Current flattening in software and hardware for security applications. In *CODES+ISSS*, pages 218–223, 2004.

[33] G. C. Necula. Proof-Carrying Code. In *Conference Record of POPL '97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 106–119, Paris, France, Jan 1997.

[34] T. Popp and S. Mangard. Masked Dual-Rail Pre-Charge Logic: DPA-Resistance without Routing Constraints. In J. R. Rao and B. Sunar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2005, 7th International Workshop, Edinburgh, Scotland, August 29 - September 1, 2005, Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 172–186. Springer, 2005.

[35] R. Ragel and S. Parameswaran. IMPRES: Integrated Monitoring for Processor REliability and Security. In *Proceedings of the Design and Automation Conference 2006 (DAC'06)*, pages 502–505, San Fransisco, California, USA, November 2006. ACM Press.

[36] R. G. Ragel. *Architectural support for security and reliability in embedded processors*. PhD thesis, University of New South Wales, 2006.

[37] H. Saputra, N. Vijaykrishnan, M. Kandemir, M. J. Irwin, R. Brooks, S. Kim, and W. Zhang. Masking the energy behavior of des encryption. *date*, 01:10084, 2003.

[38] K. Tanimura and N. Dutt. ExCCel: Exploration of complementary cells for efficient DPA attack resistivity. In *HOST*, pages 52–55, 2010.

[39] K. Tanimura and N. Dutt. HDRL: Homogeneous Dual-Rail Logic for DPA Attack Resistive Secure Circuit Design. *IEEE Embedded Systems Letters*, 4(3):57–60, 2012.

[40] K. Tanimura and N. D. Dutt. LRCG: latch-based random clock-gating for preventing power analysis side-channel attacks. In *CODES+ISSS*, pages 453–462, 2012.

[41] K. Tiri and I. Verbauwhede. A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation. In *DATE '04: Proceedings of the conference on Design, automation and test in Europe*, page 10246, Washington, DC, USA, 2004. IEEE Computer Society.

[42] K. Tiri and I. Verbauwhede. A Digital Design Flow for Secure Integrated Circuits. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 25(7):1197–1208, 2006.

[43] N. Veyrat-Charvillon, M. Medwed, S. Kerckhof, and F.-X. Standaert. Shuffling against Side-Channel Attacks: A Comprehensive Study with Cautionary Note. In *ASIACRYPT*, pages 740–757, 2012.

[44] D. Wagner, J. S. Foster, E. A. Brewer, and A. Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *Network and Distributed System Security Symposium*, pages 3–17, San Diego, CA, February 2000.

[45] S. H. Weingart. Physical security devices for computer subsystems: A survey of attacks and defences. In Çetin Kaya Koç and C. Paar, editors, *CHES*, volume 1965 of *Lecture Notes in Computer Science*, pages 302–317. Springer, 2000.

[46] W. Wolf. Multimedia applications of multiprocessor systems-on-chips. In *DATE '05: Proceedings of the conference on Design, Automation and Test in Europe*, pages 86–89, Washington, DC, USA, 2005. IEEE Computer Society.

[47] J. Xu, Z. Kalbarczyk, S. Patel, and R. Iyer. Architecture support for defending against buffer overflow attacks. In *EASY-2 Workshop*, October 2002.

[48] Y. Zhou and D. Feng. Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing. Cryptology ePrint Archive, Report 2005/388, 2005.